Biometrika (2016), **xx**, x, *pp*. 1–12
© 2007 Biometrika Trust
Printed in Great Britain

arXiv:1510.05014v2 [math.ST] 3 Mar 2016

Partial least squares for dependent data

BY MARCO SINGER, TATYANA KRIVOBOKOVA, AXEL MUNK

Institute for Mathematical Stochastics, Georg-August-Universität Göttingen, Goldschmidtstr. 7, 37077 Göttingen, Germany

msinger@gwdg.de tkrivob@uni-goettingen.de munk@math.uni-goettingen.de

AND BERT DE GROOT

Max Planck Institute for Biophysical Chemistry, Am Fassberg 11, 37077 Göttingen, Germany

bgroot@gwdg.de

SUMMARY

We consider the partial least squares algorithm for dependent data and study the consequences of ignoring the dependence both theoretically and numerically. Ignoring non-stationary dependence structures can lead to inconsistent estimation, but a simple modification leads to consistent estimation. A protein dynamics example illustrates the superior predictive power of the method.

Some key words: Dependent data, Latent variable model, Non-stationary process, Partial least squares, Protein dynamics

1. INTRODUCTION

The partial least squares algorithm introduced by Wold (1966) is a powerful regularized regression tool. It is an iterative technique, which is, unlike most similar methods, non-linear in the response variable. Consider a linear regression model

$$y = X\beta + \varepsilon, \tag{1}$$

where $y \in \mathbb{R}^n$, $X \in \mathbb{R}^{n\times k}$, $\beta \in \mathbb{R}^k$ and the error term $\varepsilon \in \mathbb{R}^n$ is a vector of n independent and identically distributed random variables. To estimate the unknown coefficients β with partial least squares, a base of $i \leq k$ weight vectors $\widehat{w}_1, \ldots, \widehat{w}_i$ is iteratively constructed. First, the data are centered, i.e., y and the columns of X are transformed to have mean zero. Then the first vector $\widehat{w}_1$ is obtained by maximizing the empirical covariance between Xw and y in $w \in \mathbb{R}^k$, subject to $\|w\| = 1$. Afterwards, the data are projected into the space orthogonal to $X\widehat{w}_1$ and the procedure is iterated. The ith partial least squares estimator $\widehat{\beta}_i$ for β is obtained by performing a least squares regression of y on X, constrained to the subspace spanned by the columns of $\widehat{W}_i = (\widehat{w}_1, \ldots, \widehat{w}_i)$. Helland (1988) summarizes the partial least squares iterations in two steps via

$$\begin{aligned} \widehat{w}_{i+1} &= b - A\widehat{\beta}_i, \quad \widehat{\beta}_0 = 0, \\ \widehat{\beta}_i &= \widehat{W}_i(\widehat{W}_i^{\mathrm{T}} A \widehat{W}_i)^{-1}\widehat{W}_i^{\mathrm{T}} b, \end{aligned} \tag{2}$$

with $b = n^{-1}X^{\mathrm{T}}y$ and $A = n^{-1}X^{\mathrm{T}}X$, under the assumption that $(\widehat{W}_i^{\mathrm{T}} A \widehat{W}_i)^{-1}$ exists. The regularisation is achieved by early stopping, that is, by taking $i \leq k$.

Alternatively, $\widehat{\beta}_i$ can be defined using the fact that $\widehat{w}_i \in \mathcal{K}_i(A,b)$, where $\mathcal{K}_i(A,b)$ is a Krylov space, that is, a space spanned by $\left\{A^{j-1}b\right\}_{j=1}^{i}$ (Helland, 1988). Then, one can define partial least squares estimators as $\widehat{\beta}_i = \arg\min_{\beta\in\mathcal{K}_i(A,b)}(y-X\beta)^{\mathrm{T}}(y-X\beta)$. There is also a direct correspondence between partial least squares and the conjugate gradient method with early stopping for the solution of $A\beta = b$.

Frank & Friedman (1993) and Farkas & Héberger (2005) find the partial least squares algorithm to be competitive with regularized regression techniques, such as principal component regression, lasso or ridge regression, in terms of the mean squared prediction error. Also, the optimal number of partial least squares base components is often much lower than that of principal components regression, as found in Almøy (1996).

Partial least squares regression has a long and successful history in various application areas, see e.g., Hulland (1999), Lobaugh et al. (2001), Nguyen & Rocke (2002). However, the statistical properties of this algorithm have been little studied, perhaps because of the non-linearity of partial least squares estimators in the response variable. Some attempts to understand properties of partial least squares can be found in Höskuldsson (1988), Phatak & de Hoog (2002) and Krämer (2007). Their almost sure convergence was established by Naik & Tsai (2000). For kernel partial least squares, Blanchard & Krämer (2010a) obtained convergence in probability results by early stopping. For the closely linked kernel conjugate gradient algorithm, Blanchard & Krämer (2010b) established order-optimal convergence rates dependent on the regularity of the target function. Delaigle & Hall (2012) compared theoretically the population and sample properties of the partial least squares algorithm for functional data.

Regression techniques typically assume independence of responses, but this is often violated, for example, if the data are observed over time or at dependent spatial locations. We are not aware of any treatment of the partial least squares algorithm for dependent observations. In this work we propose a modification of partial least squares to deal with dependent observations and study the theoretical properties of partial least squares estimators under general dependence in the data. In particular, we quantify the influence of ignored dependence.

Throughout the paper we denote by $\|\cdot\|_{\mathcal{L}}$ the spectral and by $\|\cdot\|$ the Frobenius norm for matrices, $\|\cdot\|$ also denotes the Euclidean norm for vectors.

2. PARTIAL LEAST SQUARES UNDER DEPENDENCE

2·1. *Latent variable model*

In many applications the standard linear model (1) is too restrictive. For example, if a covariate that is relevant for the response cannot be observed or measured directly, so-called latent variable or structural equation models are considered (Skrondal & Rabe-Hesketh, 2006): it is assumed that X and y are linked by $l \leq k$ latent vectors and the remaining vectors in the k-dimensional column space of X do not contribute to y. This can be interpreted as if the latent components are of interest, but only X, which contains some unknown nuisance information, can be measured. Such models are relevant in modelling of chemical (Wold et al., 2001), economic (Hahn et al., 2002) and social data (Goldberger, 1972).

We consider a latent variable model with the covariates X and response y connected via a matrix of latent variables N,

$$\begin{aligned} X &= V(NP^{\mathrm{T}} + \eta_1 F), \\ y &= V(Nq + \eta_2 f), \end{aligned} \tag{3}$$

where $N$and F are an $n \times l$-dimensional and an $n \times k$-dimensional random matrix, respectively, and f is an n-dimensional random vector. The random elements N, F, f can have different distributions, but are independent of each other, with all entries being independent and identically distributed with expectation zero and unit variance. The matrix $P \in \mathbb{R}^{k \times l}$ and vector $q \in \mathbb{R}^l$ are deterministic and unknown, along with the real-valued parameters $\eta_1, \eta_2 \geq 0$. We assume that $n \geq k \geq l$ and that $\operatorname{rank}(N) = \operatorname{rank}(P) = l$, $\operatorname{rank}(F) = k$ almost surely.

The matrix $V \in \mathbb{R}^{n \times n}$ is a deterministic symmetric matrix, such that V^2 is a positive definite covariance matrix. If $V \neq I_n$, then X in model (3) can be seen as the matrix form of a k-dimensional time series $\{X_t\}_{t=1}^n$ ($X_t \in \mathbb{R}^k$) and y can be seen as a real-valued time series $\{y_t\}_{t=1}^n$. The covariance matrix V^2 determines the dependence between observations, which might be non-stationary. We will call V^2 the temporal covariance matrix of X and define $\Sigma^2 = PP^{\mathrm{T}} + \eta_1^2 I_k$. Setting $l = k$, $\eta_1 = 0$ reduces model (3) to standard linear regression with dependent observations.

The latent variables N connect X to y, whereas F can be considered as noise, thus giving a model where not all directions in the column space of X are important for the prediction of y. The representation (3) highlights practical settings where the partial least squares algorithm is expected to outperform principal component regression and similar techniques. In particular, if the covariance of $\eta_1 F$ dominates that of NP^{T}, then the first principal components will be largely uncorrelated to y. In contrast, the first partial least squares basis components should by definition be able to recover relevant latent components.

The partial least squares algorithm is run as described in Section 1 with matrix X and vector y defined in (3). If $\eta_1 = 0$, then model (1) is correctly specified with $q = P^{\mathrm{T}}\beta$ and the partial least squares estimator (2) estimates β. If $\eta_1 > 0$, then model (1) is misspecified and $\beta(\eta_1) = \Sigma^{-2}Pq$ is rather estimated. Note that $\beta(0) = \beta$.

In the standard partial least squares algorithm it is assumed that $V = I_n$. In the subsequent sections we aim to quantify the influence of $V \neq I_n$, which is ignored in the algorithm.

2·2. *Population and sample partial least squares*

The population partial least squares algorithm for independent observations was first introduced by Helland (1990). Under model (3) we modify the definition of the population partial least squares basis vectors as

$$w_i = \arg \max_{\substack{w \in \mathbb{R}^k \\ \|w\|=1}} \frac{1}{n^2} \sum_{t,s=1}^{n} \operatorname{cov}(y_t - X_t^{\mathrm{T}}\beta_{i-1}, X_s^{\mathrm{T}} w), \quad \beta_0 = 0,$$

where $\beta_i \in \mathbb{R}^k$ are the population partial least squares regression coefficients. The average covariances over observations are taken, since the data are neither independent nor identically distributed if $V^2 \neq I_n$. Solving this optimization problem implies that the basis vectors $w_1, \ldots, w_i$ span the Krylov space $\mathcal{K}_i(\Sigma^2, Pq)$: see the Supplementary Material. In particular, under model (3), the Krylov space in the population turns out to be independent of the temporal covariance V^2 for all $n \in \mathbb{N}$.

For a given Krylov space, the population partial least squares coefficients are obtained as

$$\beta_i = \arg \min_{\beta \in \mathcal{K}_i(\Sigma^2, Pq)} E\left\{ \frac{1}{n} \sum_{t=1}^{n} (y_t - X_t^{\mathrm{T}}\beta)^2 \right\}.$$

It is easy to see that the solution to this problem is

$$\beta_i = K_i \left(K_i^{\mathrm{T}} \Sigma^2 K_i \right)^{-1} K_i^{\mathrm{T}} Pq, \quad K_i = (Pq, \Sigma^2 Pq, \ldots, \Sigma^{2(i-1)} Pq),$$

which is independent of V^2 for all $n \in \mathbb{N}$.

To obtain the sample partial least squares estimators $\widehat{\beta}_i$, Σ^2 and Pq are replaced by estimators. In the standard partial least squares algorithm, under independence of observations, Σ^2 and Pq are estimated by unbiased estimators $n^{-1}X^{\mathrm{T}}X$ and $n^{-1}X^{\mathrm{T}}y$, respectively. However, if the observations are dependent, such naive estimators can lead to L_2-inconsistent estimation, as the following theorem shows.

THEOREM 1. *Let the model (3) hold and the fourth moments of $N_{1,1}$, $F_{1,1}$ exist. Define $A = \|V\|^{-2}X^{\mathrm{T}}X, \quad b = \|V\|^{-2}X^{\mathrm{T}}y$. Then*

$$E\left(\left\|\Sigma^2 - A\right\|^2\right) = \frac{\|V^2\|^2}{\|V\|^4}\left(C_A + \sum_{t=1}^{n}\frac{\|V_t\|^4}{\|V^2\|^2}c_A\right)$$
$$E\left(\|Pq - b\|^2\right) = \frac{\|V^2\|^2}{\|V\|^4}\left(C_b + \sum_{t=1}^{n}\frac{\|V_t\|^4}{\|V^2\|^2}c_b\right),$$

where

$$C_A = \|P\|^4 + \|P^{\mathrm{T}}P\|^2 + 4\eta_1^2\|P\|^2 + \eta_1^4 k(1+k)$$
$$c_A = \left\{E\left(N_{1,1}^4\right) - 3\right\}\sum_{i=1}^{l}\|P_i\|^4 + \left\{E\left(F_{1,1}^4\right) - 3\right\}\eta_1^4 k$$
$$C_b = \|Pq\|^2 + \|P\|^2\|q\|^2 + \eta_1^2 k\|q\|^2 + \eta_1^2\eta_2^2 k + \eta_2^2\|P\|^2$$
$$c_b = \left\{E\left(N_{1,1}^4\right) - 3\right\}\sum_{i=1}^{l}\|P_i\|^2 q_i^2$$

and V_t denotes the t-th column of matrix V.

The scaling factors in A and b have no influence on the sample partial least squares estimators in (2), so that replacing n^{-1} with $\|V\|^{-2}$ does not affect the algorithm and both A and b are unbiased estimators for Σ^2 and Pq, respectively.

If $E(N_{1,1}^4) = E(F_{1,1}^4) = 3$, then constants c_A and c_b vanish, simplifying expressions for the mean squared error of A and b. This is satisfied, for example, for the standard normal distribution. Thus, these terms can be interpreted as a penalization for non-normality.

Finally, $\sum_{t=1}^{n}\|V_t\|^4 \leq \sum_{t,s=1}^{n}(V_t^{\mathrm{T}}V_s)^2 = \left\|V^2\right\|^2$ implies that the convergence rate of both estimators is driven by the ratio of Frobenius norms $\|V\|^{-2}\|V^2\|$. In particular, if $\|V\|^{-2}\|V^2\|$ converges to zero, then the elements of the population Krylov space Σ^2 and Pq can be estimated consistently. This is the case, for example, for independent observations with $V = I_n$, since $\|I_n^2\| = \|I_n\| = n^{1/2}$. However, if $\|V\|^{-2}\|V^2\|$ fails to converge to zero, ignoring the temporal dependence V^2 may lead to inconsistent estimation.

3. PROPERTIES OF PARTIAL LEAST SQUARES ESTIMATORS UNDER DEPENDENCE

3·1. *Concentration inequality for partial least squares estimators*

In this section we apply techniques of Blanchard & Krämer (2010b), who derived convergence rates of the kernel conjugate gradient algorithm, which is closely related to kernel partial least squares. Both algorithms approximate the solution on Krylov subspaces, but employ different norms. In particular, Blanchard & Krämer (2010b) have shown that if the conjugate gradient

algorithm is stopped early, the convergence in probability of the kernel conjugate gradient estimator to the true regression function can be obtained for bounded kernels. Moreover, the convergence is order-optimal, depending on the regularity of the target function. These results hold for independent identically distributed observations.

We avoid the nonparametric setting of Blanchard & Krämer (2010b) and study a standard linear partial least squares algorithm with a fixed dimension k of the regression space. We allow the observations to be dependent, and, instead of a bounded kernel, consider unbounded random variables with moment conditions. In this setting we derive concentration inequalities for partial least squares estimators that allow us to quantify the influence of the temporal covariance.

Regularization of the partial least squares solution is achieved by early stopping, which is characterized by the discrepancy principle, i.e., we stop at the first index $0 < a_0 \leq a$ such that

$$\left\| A^{1/2}\widehat{\beta}_{a_0} - A^{-1/2}b \right\| \leq \tau(\delta\|\widehat{\beta}_{a_0}\| + \epsilon), \tag{4}$$

for $\delta, \epsilon > 0$ defined in Theorem 2, and some $\tau \geq 1$. Here a denotes the maximal dimension of the sample Krylov space $\mathcal{K}_i(A, b)$ and almost surely equals $a = l + (k - l)\mathbb{I}(\eta_1 > 0)$, where $\mathbb{I}(\cdot)$ denotes an indicator function. For technical reasons we stop at $a^* = a_0 - 1$ if $p_{a_0}(0) \geq \zeta\delta^{-1}$, where p_i is a polynomial of degree $i - 1$ with $p_i(A)b = \widehat{\beta}_i$ and $\zeta < \tau^{-1}$. The existence of such polynomials was proved by Phatak & de Hoog (2002). If (4) never holds, $a^* = a$ is taken. With this stopping index we get the following concentration inequality.

THEOREM 2. *Assume that model (3) with $\eta_1 > 0$ holds and that the fourth moments of $N_{1,1}$, $F_{1,1}$ exist. Furthermore, a^* satisfies (4) with $\tau \geq 1$, $\zeta < \tau^{-1}$. For $\nu \in (0, 1]$ let $\delta = \nu^{-1/2}\|V\|^{-2}\|V^2\|C_\delta$ and $\epsilon = \nu^{-1/2}\|V\|^{-2}\|V^2\|C_\epsilon$, such that $\delta, \epsilon \to 0$, where*

$$C_\delta = (2C_A + 2c_A)^{1/2}, \quad C_\epsilon = (2C_b + 2c_b)^{1/2},$$

with C_A, c_A, C_b and c_b given in Theorem 1. Then with a probability at least $1 - \nu$,

$$\left\| \widehat{\beta}_{a^*} - \beta(\eta_1) \right\| \leq \frac{\|V^2\|}{\|V\|^2} \left\{ c_1(\nu) + \frac{\|V^2\|}{\|V\|^2} c_2(\nu) \right\}, \tag{5}$$

where

$$c_1(\nu) = \nu^{-1/2} c(\tau, \zeta)\|\Sigma^{-1}\|_{\mathcal{L}} \left(C_\epsilon + \|\Sigma\|_{\mathcal{L}}\|\Sigma^{-3}Pq\|C_\delta \right)$$
$$c_2(\nu) = \nu^{-1} c(\tau, \zeta)\|\Sigma^{-1}\|_{\mathcal{L}} \left(C_\epsilon C_\delta + \|\Sigma^{-3}Pq\|C_\delta^2 \right),$$

for some constant $c(\tau, \zeta)$ that asymptotically depends only on τ and ζ.

If $N_{1,1}$,$F_{1,1}$, $f_1 \sim \mathcal{N}(0, 1)$, then the expressions for C_δ and C_ϵ are simplified and the scaling factor of $c_1(\nu)$ and $c_2(\nu)$ can be improved from $\nu^{-1/2}$ to $\log(2/\nu)$, which is achieved by using an exponential inequality proved in Theorem 3.3.4 of Yurinsky (1995).

Theorem 2 states that the convergence rate of the optimally stopped partial least squares estimator $\widehat{\beta}_{a^*}$ to the true parameter $\beta(\eta_1)$ is driven by the ratio of the Frobenius norms of V^2 and V, similar to the results of Theorem 1. In particular, if the data are independent with $V = I_n$ then $\widehat{\beta}_{a^*}$ is square-root consistent. In this case $c_2(\nu)$ is asymptotically negligible. Note that the theorem excludes the case that $\|V\|^{-2}\|V^2\|$ does not converge to 0.

3·2. *Properties of $\widehat{\beta}_1$ under dependence*

Non-linearity in the response variable of $\widehat{\beta}_i$ hinders its standard statistical analysis, as no closed-form expression for the mean square error of $\widehat{\beta}_i$ is available and concentration inequalities

similar to (5) are the only results on the convergence rates of partial least squares estimators, to the best of our knowledge. However, if the ratio of $\|V^2\|$ and $\|V\|^2$ does not converge to zero, Theorem 2 does not hold.

In this section we study the first partial least squares estimator $\widehat{\beta}_1$, for several reasons. First, the explicit expression for its mean square error can be derived. Second, if there is only one latent component that links X and y, i.e., $l = 1$ in (3), then consistent estimation of β_1 is crucial. Finally, $\widehat{\beta}_1$ is collinear to the direction of the maximal covariance between X and y given by $\widehat{w}_1$, which is important for the interpretation of the partial least squares model in applications, see Krivobokova et al. (2012). The next theorem gives conditions under which $\widehat{\beta}_1$ is an inconsistent estimator of β_1.

THEOREM 3. *Assume that model (3) holds, $k > 1$ and eighth moments of $N_{1,1}$, $F_{1,1}$, f_1 exist. Furthermore, suppose that the ratio $\|V\|^{-2}\|V^2\|$ does not converge to zero as $n \to \infty$. Then, for either $l > 1$, $\eta_1 \geq 0$ or $l = 1$, $\eta_1 > 0$, $\widehat{\beta}_1$ is an inconsistent estimator for β_1.*

The case $l = 1$, $\eta_1 = 0$ not treated in Theorem 3 corresponds to the standard linear regression model with a single covariate, so the partial least squares estimator coincides with the ordinary least squares estimator, see Helland (1988).

Hence, if there is only one latent component in the model, i.e., $l = 1$, $\eta_1 > 0$, and $\|V\|^{-2}\|V^2\|$ does not converge to zero, then $\beta(\eta_1)$, which in this case equals β_1, cannot be estimated consistently with a standard partial least squares algorithm.

3·3. *Examples of dependence structures*

In all previous theorems the ratio $\|V^2\|\|V\|^{-2}$ plays a crucial role. In this section some special covariance matrices V^2 are studied in order to understand its behaviour. Stationary processes considered in this section are assumed to have expectation zero and to decay exponentially, i.e., for for $c, \rho > 0$ and $\gamma(0) > 0$,

$$|\gamma(t)| \leq \gamma(0) c \exp(-\rho t), \quad t \in \mathbb{N}, \tag{6}$$

with $\gamma : \mathbb{Z} \to \mathbb{R}$ being the autocovariance function of the process.

Subsequently, $f(n) \sim g(n)$ denotes $c_1 \leq f(n)/g(n) \leq c_2$, for n large, $0 < c_1 < c_2$ and $f, g : \mathbb{N} \to \mathbb{R}$.

THEOREM 4. *Let $\left[V^2\right]_{t,s} = \gamma(|t-s|)$ $(t, s = 1, \ldots, n)$ be the covariance matrix of a stationary process, such that the autocovariance function $\gamma : \mathbb{Z} \to \mathbb{R}$ satisfies (6). Then $\|V^2\| \sim n^{1/2}$ and $\|V\|^2 \sim n$.*

Hence, if V^2 in model (3) is a covariance matrix of a stationary process, then ignoring dependence of observations in the partial least squares algorithm does not affect the rate of convergence of partial least squares estimators, but might affect the constants. Examples of processes with exponentially decaying autocovariances are stationary autoregressive moving average processes.

As an example of a non-stationary process we consider first-order integrated processes. If $\{X_t\}_{t\in\mathbb{Z}}$ is stationary with autocovariance function γ satisfying (6), then $\sum_{i=1}^{t} X_i$ is an integrated process of order 1.

THEOREM 5. *Let $\{X_t\}_{t\in\mathbb{Z}}$ be a stationary process with autocovariance function γ satisfying (6). If $\gamma(t) < 0$ for some t, we assume additionally $\rho > \log(2c+1)$. Let V^2 be the covariance matrix of $\sum_{i=1}^{t} X_i$. Then $\|V\|^2 \sim n^2$ and $\|V^2\| \sim n^2$.*

The lower bound on ρ for negative $\gamma(t)$ ensures that no element on the diagonal of V^2 becomes negative, so that V^2 is a valid covariance matrix.

This theorem implies that the ratio $\|V\|^{-2}\|V^2\|$ does not converge to zero for certain integrated processes. In particular, combining this result with Theorems 1 and 3 shows that the elements of the sample Krylov space A and b, as well as $\widehat{\beta}_1$, are inconsistent, if the dependence structure of the data can be described by an integrated process satisfying the conditions of Theorem 5, e.g., an integrated autoregressive moving average process of order $(1, 1, 1)$.

4. PRACTICAL ISSUES

4·1. *Corrected partial least squares estimator*

So far we considered the standard partial least squares algorithm, showing that if certain dependences in the data are ignored, estimation is inconsistent. Hence, it is crucial to take into account the dependence structure of the data in the partial least squares estimators.

Let us define $b(S) = n^{-1}X^{\mathrm{T}}S^{-2}y$ and $A(S) = n^{-1}X^{\mathrm{T}}S^{-2}X$ for an invertible matrix $S \in \mathbb{R}^{n\times n}$. Furthermore, let $k_i(S) = A(S)^{i-1}b(S)$, $K_i(S) = [k_1(S), \ldots, k_i(S)] \in \mathbb{R}^{k\times i}$ and $\widehat{\beta}_i(S) = K_i(S)\{K_i(S)^{\mathrm{T}}A(S)K_i(S)\}^{-1}K_i(S)^{\mathrm{T}}b(S)$ $(i = 1, \ldots, k)$.

For $S = I_n$ this yields a standard partial least squares estimator. If $S = V$, the temporal dependence matrix, then $b(V)$ and $A(V)$ are square-root consistent estimators of Pq and Σ^2, respectively, with the mean squared error independent of V, which follows from Theorem 1. Hence, the resulting $\widehat{\beta}_i(V)$ is also a consistent estimator of β_i and Theorem 2 shows that $\beta(\eta_1)$ can be estimated consistently by early stopping as well. This procedure is equivalent to running the partial least squares algorithm on $V^{-1}y$ and $V^{-1}X$, that is, with the temporal dependence removed from the data.

In practice the true covariance matrix V^2 is typically unknown and is replaced by a consistent estimator $\widehat{V}^2$. We call the estimator $\widehat{\beta}_i(\widehat{V})$ the corrected partial least squares estimator. The next theorem shows that, given a consistent estimator of V^2, the population Krylov space and $\beta(\eta_1)$ can be estimated consistently.

THEOREM 6. *Let $\widehat{V}^2$ be an estimator for V^2 that is invertible for all $n \in \mathbb{N}$ and $\left\|V\widehat{V}^{-2}V - I_n\right\|_{\mathcal{L}} = O_p(r_n)$, where r_n is some sequence of positive numbers such that $r_n \to 0$ as $n \to \infty$. Then*

$$\|A(\widehat{V}) - \Sigma^2\|_{\mathcal{L}} = O_p(r_n), \quad \|b(\widehat{V}) - Pq\| = O_p(r_n).$$

Moreover, with probability at least $1 - \nu$, $\nu \in (0, 1]$

$$\|\widehat{\beta}_{a^*}(\widehat{V}) - \beta(\eta_1)\| = O(r_n),$$

where the definition of a^* in (4) is updated by replacing A, b and $\widehat{\beta}_i$ by $A(\widehat{V})$, $b(\widehat{V})$ and $\widehat{\beta}_i(\widehat{V})$, respectively.

Theorem 6 states that if a consistent estimator of the covariance matrix V^2 is available, then the elements of the population Krylov space A, b, as well as the coefficient $\beta(\eta_1)$, can be consistently estimated by $A(\widehat{V})$, $b(\widehat{V})$ and $\widehat{\beta}_{a^*}(\widehat{V})$. The convergence rate of these estimators is not faster than that of $\widehat{V}^2$. For example, if the temporal dependence in the data follows some parametric model, then parametric rates of $n^{-1/2}$ are also achieved for $A(\widehat{V})$, $b(\widehat{V})$ and $\widehat{\beta}_{a^*}(\widehat{V})$. Estimation of V^2 by some nonparametric methods, e.g., with a banding or tapering approach, leads to a slower convergence rates: see Bickel & Levina (2008) or Wu & Xiao (2012). Similar results are well-known in the context of linear regression. For example, Theorem 5.7.1 in Fuller (1996) shows

that the convergence rate of feasible generalized least squares estimators is the same as that of the estimator for the covariance matrix of the regression error.

4·2. *Estimation of covariance matrices*

To obtain the corrected partial least squares estimator, some consistent estimator of V^2 based on a single realisation of the process is necessary. In model (3) the dependence structure over the observations of X is the same as that of y and V can be estimated from y alone.

If V^2 is the autocovariance matrix of a stationary process, it can be estimated both parametrically and nonparametrically. Many stationary processes can be sufficiently well approximated by an autoregressive moving average process, see Brockwell & Davis (1991), Chapter 4.4. Parameters of autoregressive moving average processes are estimated either by Yule–Walker or maximum likelihood estimators, both attaining parametric rates. Another approach is to band or taper the empirical autocovariance function of y (Bickel & Levina, 2008; Wu & Pourahmadi, 2009; Wu & Xiao, 2012). These nonparametric estimators are very flexible, but are computationally intensive and have slower convergence rates.

If y is an integrated processes of order one, then V^2 can easily be derived from the covariance matrix estimator of the corresponding stationary process.

5. SIMULATIONS

To verify small sample performance of the partial least squares algorithm under dependence we consider the following simulation setting. To illustrate consistency we choose three sample sizes $n \in \{250, 500, 2000\}$. In the latent variable model (3) we set $k = 20$, $l = 1, 5$ and take the elements of P to be independent identically distributed Bernoulli random variables with success probability 0.5. Elements of the vector q are $q_i = 5\, i^{-1}$ ($i = 1, \ldots, l$), in order to control the importance of the different latent variables for y. The random variables $N_{1,1}$, $F_{1,1}$ and f_1 are taken to be standard normally distributed. The parameter η_2 is chosen to get the signal to noise ratio in y to be two and η_1 is set so that the signal to noise ratio in X is 0.5. Three matrices V^2 are considered: the identity matrix, the covariance matrix of an autoregressive process of the first order with coefficient 0.9 and the covariance matrix of an autoregressive integrated moving average process of order $(1, 1, 1)$ with both parameters set to 0.9.

First, we ran the standard partial least squares algorithm on the data with the three aforementioned dependence structures to highlight the effect of the ignored dependence in the data. Next, we studied the performance of our corrected partial least squares algorithm applied to nonstationary data. Thereby, the covariance matrix of the autoregressive moving average process has been estimated parametrically, as discussed in Section 4·2. A nonparametric estimation of this covariance matrix has lead to qualitative similar results.

The boxplots in Figure 1 show the squared distance of $\widehat{\beta}_i$ and $\beta(\eta_1)$ in 500 Monte Carlo replications. Two cases are shown in one panel: the model has just one latent component and $\widehat{\beta}_1$ is considered, i.e., $l = i = 1$ and the model has five latent components and the squared distance of $\widehat{\beta}_5$ to $\beta(\eta_1)$ is studied, i.e., $l = i = 5$.

We observe that the mean squared error of $\widehat{\beta}_i$ obtained with the standard partial least squares converges to zero for autoregressive and independent data with the growing sample size. However, an autoregressive dependence in the data leads to a somewhat higher mean squared error, compare the top and bottom left panels. If the data follow an autoregressive integrated moving average process and this is ignored in the partial least squares algorithm, then the mean squared error of $\widehat{\beta}_i$ converges to some positive constant, see the top right boxplots. Taking into account

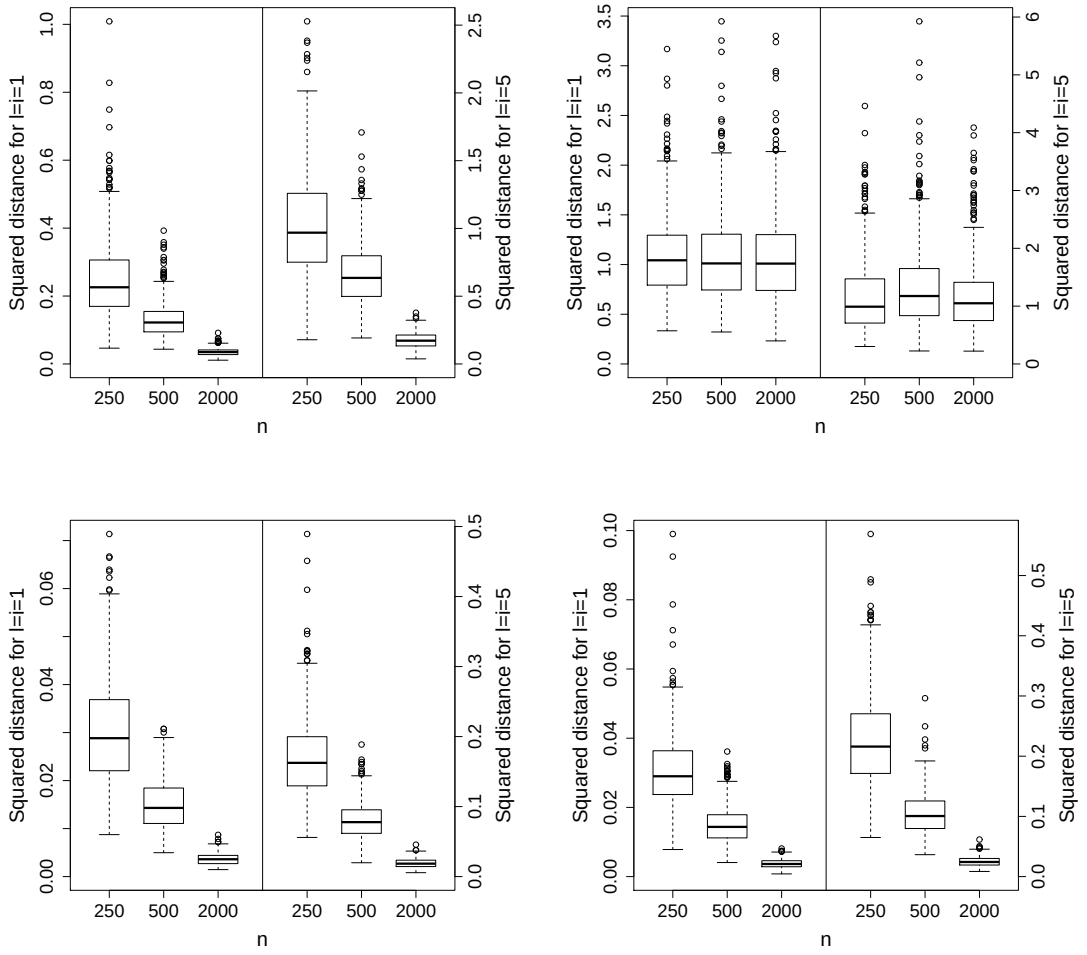

Fig. 1: Squared distance of partial least squares estimators $\widehat{\beta}_i$ and $\beta(\eta_1)$ in 500 Monte Carlo samples. First three boxplots in each panel correspond to $l = i = 1$, the latter three to $l = i = 5$. The dependence structures are: first order autoregressive (top left), autoregressive integrated moving average of order (1,1,1) (right) and independent, identically distributed (bottom left). The standard partial least squares (top and bottom left) and corrected partial least squares (bottom right) have been employed.

these non-stationary dependencies in the corrected partial least squares leads to consistent estimation, similar to the independent data case, compare the bottom left and right panels.

We conclude that if the observations are dependent, corrected partial least squares improves estimation: in case of stationary dependence the mean squared error is reduced and in case of non-stationary dependence the estimation becomes consistent.

6. Application to Protein Dynamics

Proteins fulfil their biological function through particular movements, see Henzler-Wildman

& Kern (2007), so a key step in understanding protein functions is a detailed knowledge of the underlying dynamics. Molecular dynamics simulations (de Groot et al., 1998) are routinely used to study the dynamics of biomolecular systems at atomic detail on timescales of nanoseconds to microseconds. Although in principle allowing to directly address function-dynamics relationships, analysis is frequently hampered by the large dimensionality of the protein configuration space, rendering it non-trivial to identify collective modes of motion that are directly related to a functional property of interest.

Krivobokova et al. (2012) have shown that partial least squares helps to identify a hidden relation between atom coordinates of a protein and a functional parameter of interest, yielding robust and parsimonious solutions, superior to principal component regression. In this work we look at a protein studied in the aforementioned paper: the water channel aquaporine as found in the yeast Pichia pastoris. This is a gated channel, i.e., the diameter of the opening can change, controlling the flow of water into the cell. We aim to study which collective motions of protein atoms influence the diameter y_t of the channel at time t, measured by the distance of two centres of mass of the residues of the protein which characterize the opening. For the description of the protein dynamics we use an inner model, i.e. at each point in time we calculate the Euclidean distance d of each backbone atom of the protein and a set of certain four fixed base points. We denote the $p = 739$ atoms by $A_{t,1}, \ldots, A_{t,p} \in \mathbb{R}^3$, the fixed base points by $B_1, \ldots, B_4 \in \mathbb{R}^3$ and take

$$X_t = \{d(A_{t,1}, B_1), \ldots, d(A_{t,p}, B_1), d(A_{t,1}, B_2), \ldots, d(A_{t,p}, B_4)\}^{\mathrm{T}} \in \mathbb{R}^{4p}.$$

The available timeframe has a length of 100 ns split into $n = 20\,000$ equidistant points of observation. Krivobokova et al. (2012) found that a linear relationship between X and y can be assumed. Additionally, these data seem not to contradict model (3). Taking a closer look at the data reveals that both y_t and $X_{t,i}$ ($i = 1, \ldots, 4p$) are non-stationary time series, see Figure 2. For the calculation of $\widehat{V}^2$ we used the banding approach mentioned in Section 4·2 and found the results to be very similar to a simple autoregressive integrated moving average process with parameters (3,1,1) and corresponding coefficients $(0.1094, 0.0612, 0.0367, -0.9159)$. Autore-

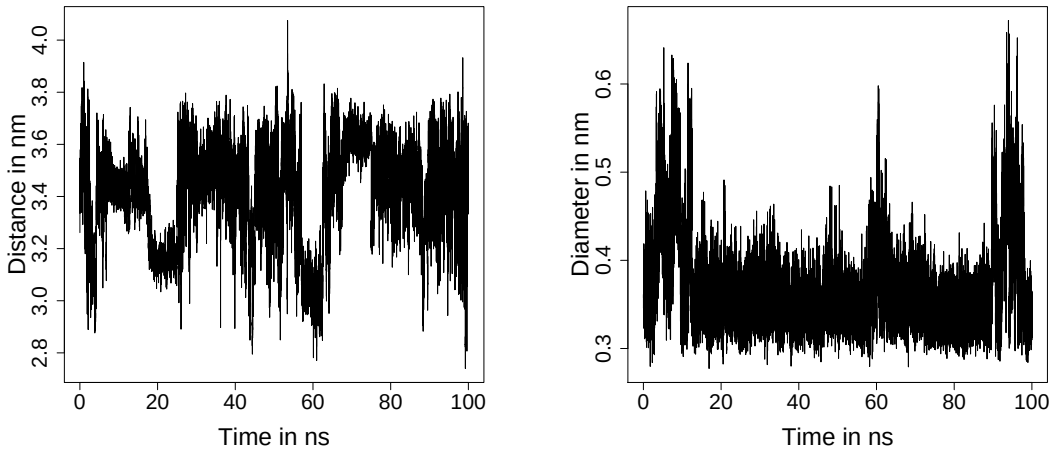

Fig. 2: Distance between the first backbone atom and the first centre of mass of aquaporine (left) and the opening diameter over time (right).

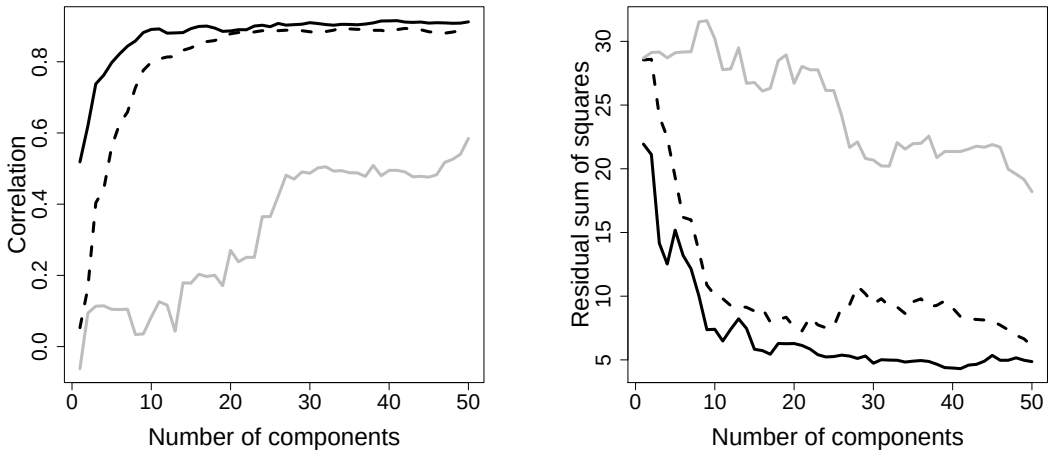

Fig. 3: Correlation (left) and residual sum of squares (right) of the predicted opening diameter and the real data on the test set. Compared methods are principal component regression (grey), corrected partial least squares (black, solid) and partial least squares (black, dashed).

gressive integrated moving average models have been employed before to study protein time series (Alakent et al., 2004).

To validate our estimators, we used the following procedure. First, the data were split into two equal parts and the models were build on the first half. Then the prediction was done on the test set consisting of the second half of the data and was compared to y_t from the test set. To measure the accuracy of the prediction we used the Pearson correlation coefficient common in the biophysics community and the residual sum of squares, both shown in Figure 3. The partial least squares estimator clearly outperforms principal components regression. The corrected partial least squares algorithm, which takes temporal dependence into account, delivers better prediction than standard partial least squares. The improvement is strongly present in the first components.

High predictive power of the first corrected partial least squares components is particularly relevant for the interpretation of the underlying protein dynamics. Krivobokova et al. (2012) established that the first partial least squares regression coefficient $\widehat{\beta}_1$ corresponds to the so-called ensemble-weighted maximally correlated mode of motion contributing most to the fluctuation in the response y. Altogether, due to the low dimensionality, corrected partial least squares greatly facilitates the interpretation of the underlying relevant dynamics, compared to partial least squares and principal component regression, where many more components are required to obtain the same predictive power.

ACKNOWLEDGEMENTS

The authors are grateful to the reviewers for their valuable comments which helped to improve the article. The support of the German Research Foundation is gratefully acknowledged.

SUPPLEMENTARY MATERIAL

Supplementary Material available at *Biometrika* online includes all technical details.

REFERENCES

ALAKENT, B., DORUKER, P. & CAMURDAN, M. (2004). Time series analysis of collective motions in proteins. *J. Chem. Phys.* **120**, 1072–1088.

ALMØY, T. (1996). A simulation study on the comparison of prediction methods when only a few components are relevant. *Comput. Statist. Data Anal.* **21**, 87–107.

BICKEL, P. & LEVINA, E. (2008). Regularized estimation of large covariance matrices. *Ann. Statist.* **36**, 199–227.

BLANCHARD, G. & KRÄMER, N. (2010a). Kernel partial least squares is universally consistent. In *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics, JMLR Workshop and Conference Proceedings*, Y. W. Teh, ed., vol. 9. JMLR, pp. 57–64.

BLANCHARD, G. & KRÄMER, N. (2010b). Optimal learning rates for kernel conjugate gradient regression. *Adv. Neural Inf. Process. Syst.* **23**, 226–234.

BROCKWELL, P. & DAVIS, R. (1991). *Time Series: Theory and Methods*. New York: Springer, 2nd ed.

DE GROOT, B. L., HAYWARD, S., VAN AALTEN, D. M. F., AMADEI, A. & BERENDSEN, H. J. C. (1998). Domain motions in bacteriophage T4 lysozyme; a comparison between molecular dynamics and crystallographic data. *Proteins* **31**, 116–127.

DELAIGLE, A. & HALL, P. (2012). Methodology and theory for partial least squares applied to functional data. *Ann. Statist.* **40**, 322–352.

FARKAS, O. & HÉBERGER, K. (2005). Comparison of ridge regression, partial least-squares, pairwise correlation, forward-and best subset selection methods for prediction of retention indices for aliphatic alcohols. *J. Chem. Inf. Model.* **45**, 339–346.

FRANK, I. & FRIEDMAN, J. H. (1993). A statistical view of some chemometrics regression tools. *Technometrics* **35**, 109–135.

FULLER, W. (1996). *Introduction to Statistical Time Series*. New York: Wiley, 2nd ed.

GOLDBERGER, A. (1972). Structural equation methods in the social sciences. *Econometrica* **40**, 979–1001.

HAHN, C., JOHNSON, M., HERRMANN, A. & HUBER, F. (2002). Capturing customer heterogeneity using a finite mixture PLS approach. *Schmalenbach Business Review* **54**, 243–269.

HELLAND, I. S. (1988). On the structure of partial least squares regression. *Comm. Statist. Simulation Comput.* **17**, 581–607.

HELLAND, I. S. (1990). Partial least squares regression and statistical models. *Scand. J. Statist.* **17**, 97–114.

HENZLER-WILDMAN, K. & KERN, D. (2007). Dynamic personalities of proteins. *Nature* **450**, 964–972.

HÖSKULDSSON, A. (1988). PLS regression methods. *J. Chemometr.* **2**, 211–228.

HULLAND, J. (1999). Use of partial least squares (PLS) in strategic management research: A review of four recent studies. *Strateg. Manage. J.* **20**, 195–204.

KRÄMER, N. (2007). An overview on the shrinkage properties of partial least squares regression. *Comput. Statist.* **22**, 249–273.

KRIVOBOKOVA, T., BRIONES, R., HUB, J., MUNK, A. & DE GROOT, B. (2012). Partial least squares functional mode analysis: application to the membrane proteins AQP1, Aqy1 and CLC-ec1. *Biophys. J.* **103**, 786–796.

LOBAUGH, N., WEST, R. & MCINTOSH, A. (2001). Spatiotemporal analysis of experimental differences in event-related potential data with partial least squares. *Psychophysiology* **38**, 517–530.

NAIK, P. & TSAI, C.-L. (2000). Partial least squares estimator for single-index models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **62**, 763–771.

NGUYEN, D. & ROCKE, D. (2002). Tumor classification by partial least squares using microarray gene expression data. *Bioinformatics* **18**, 39–50.

PHATAK, A. & DE HOOG, F. (2002). Exploiting the connection between PLS, Lanczos methods and conjugate gradients: alternative proofs of some properties of PLS. *J. Chemometr.* **16**, 361–367.

SKRONDAL, A. & RABE-HESKETH, S. (2006). Latent variable modelling: A survey. *Scand. J. Statist.* **34**, 712–745.

WOLD, H. (1966). Nonlinear estimation by iterative least squares procedure. In *Research papers in statistics: Festschrift for J. Neyman*, F. N. David, ed. Wiley, pp. 411–444.

WOLD, S., SJØSTRØMA, M. & ERIKSSON, L. (2001). PLS-regression: a basic tool of chemometrics. *Chemometr. Intell. Lab.* **58**, 109–130.

WU, W. & POURAHMADI, M. (2009). Banding sample autocovariance matrices of stationary processes. *Statist. Sinica* **19**, 1755–1768.

WU, W. & XIAO, H. (2012). Covariance matrix estimation in time series. *Ann. Statist.* **40**, 466–493.

YURINSKY, V. (1995). *Sums and Gaussian Vectors*. Berlin: Springer, 1st ed.

arXiv:1510.05014v2 [math.ST] 3 Mar 2016

Biometrika (2016), **xx**, x, *pp.* 1–7
© 2007 Biometrika Trust
Printed in Great Britain

Supplementary material for Partial least squares for dependent data

BY MARCO SINGER, TATYANA KRIVOBOKOVA, AXEL MUNK

Institute for Mathematical Stochastics, Goalgorithm2e.styldschmidtstr. 7, 37077 Göttingen, Germany

msinger@gwdg.de tkrivob@uni-goettingen.de munk@math.uni-goettingen.de

AND BERT DE GROOT

Max Planck Institute for Biophysical Chemistry, Am Fassberg 11, 37077 Göttingen, Germany

bgroot@gwdg.de

S1. PROOFS

S1·1. *Derivation of the population partial least squares components*

Let denote $K_i \in \mathbb{R}^{k\times i}$ the matrix representation of a base for $\mathcal{K}_i(\Sigma^2, Pq)$. Then

$$\sum_{t=1}^{n} E\left(y_t - X_t^{\mathrm{T}} K_i \alpha\right)^2 = \sum_{t=1}^{n} [V^2]_{t,t} \left(\|q\|^2 + \eta_2^2 - 2\alpha^{\mathrm{T}} K_i^{\mathrm{T}} Pq + \alpha^{\mathrm{T}} K_i^{\mathrm{T}} \Sigma^2 K_i \alpha\right).$$

Minimizing this expression with respect to $\alpha \in \mathbb{R}^i$ gives $K_i^{\mathrm{T}}\Sigma^2 K_i \alpha = K_i Pq$. Since the matrix $K_i^{\mathrm{T}}\Sigma^2 K_i$ is invertible, we get the least squares fit β_i in Section 2.

Assume now that the first $i < a$ partial least squares base vectors $w_1, \ldots, w_i$ have been calculated and consider for $\lambda \in \mathbb{R}$ the Lagrange function

$$\sum_{t,s=1}^{n} \operatorname{cov}\left(y_t - X_t^{\mathrm{T}}\beta_i, X_s^{\mathrm{T}} w\right) - \lambda(\|w\|^2 - 1) = w^{\mathrm{T}}\left(Pq - \Sigma^2\beta_i\right) \sum_{t,s=1}^{n} [V^2]_{t,s} - \lambda(\|w\|^2 - 1).$$

Maximizing with respect to w yields

$$w_{i+1} = (2\lambda)^{-1}\left(Pq - \Sigma^2\beta_i\right) \sum_{t,s=1}^{n} [V^2]_{t,s} \propto Pq - \Sigma^2\beta_i.$$

Since $\beta_i \in \mathcal{K}_i(\Sigma^2, Pq)$, we get $w_{i+1} \in \mathcal{K}_{i+1}(\Sigma^2, Pq)$ and w_{i+1} is orthogonal to $w_1, \ldots, w_i$.

S1·2. *Proof of Theorem 1*

First consider

$$\begin{aligned} E\left(\|b - Pq\|^2\right) =& E\left[\left\| \frac{1}{\|V\|^2}\left\{(PN^{\mathrm{T}} + \eta_1 F^{\mathrm{T}})V^2 Nq + \eta_2 (PN^{\mathrm{T}} + \eta_1 F^{\mathrm{T}})V^2 f\right\} - Pq\right\|^2\right] \\ =& \left\{ E\left(\left\| \frac{1}{\|V\|^2} PN^{\mathrm{T}} V^2 Nq - Pq\right\|^2\right) + \frac{\eta_2^2}{\|V\|^4} E\left(\left\|PN^{\mathrm{T}} V^2 f\right\|^2\right)\right\} \\ &+ \frac{\eta_1^2}{\|V\|^4}\left\{E\left(\left\|F^{\mathrm{T}} V^2 Nq\right\|^2\right) + \eta_2^2 E\left(\left\|F^{\mathrm{T}} V^2 f\right\|^2\right)\right\} = S_1 + S_2, \end{aligned}$$

due to the independence of N, F and f. It is easy to see that

$$S_2 = \frac{\|V^2\|^2}{\|V\|^4}\eta_1^2 k\left(\|q\|^2 + \eta_2^2\right).$$

Furthermore, with notation $A_0 = N^{\mathrm{T}}V^2N$ we get

$$S_1 = \frac{1}{\|V\|^4}E\left(q^{\mathrm{T}}A_0P^{\mathrm{T}}PA_0q\right) - \|Pq\|^2 + \frac{\eta_2^2}{\|V\|^4}E\left(\left\|PN^{\mathrm{T}}V^2f\right\|^2\right).$$

Consider now $E\left(q^{\mathrm{T}}A_0P^{\mathrm{T}}PA_0q\right)$ as a quadratic form with respect to the matrix $P^{\mathrm{T}}P$. Denote $\kappa = E\left(N_{1,1}^4\right) - 3$. First, $E\left(A_0q\right) = E\left(N^{\mathrm{T}}V^2Nq\right) = \|V\|^2q$ and

$$\begin{aligned}
\mathrm{var}(A_0q) &= \left[\sum_{a,b=1}^{l} q_aq_b \sum_{t,s,u,v=1}^{n} V_u^{\mathrm{T}}V_sV_t^{\mathrm{T}}V_vE(N_{s,i}N_{u,a}N_{t,j}N_{v,b})\right]_{i,j=1}^{l} - \|V\|^4qq^{\mathrm{T}} \\
&= \left[q_iq_j\|V\|^4 + \left(q_iq_j + \delta_{i,j}\|q\|^2\right)\|V^2\|^2 + \kappa\sum_{t=1}^{n}\|V_t\|^4\delta_{i,j}q_i^2\right]_{i,j=1}^{l} - \|V\|^4qq^{\mathrm{T}} \\
&= \|V^2\|^2\left(qq^{\mathrm{T}} + \|q\|^2I_l\right) + \kappa\sum_{t=1}^{n}\|V_t\|^4\mathrm{diag}\left(q_1^2,\ldots,q_l^2\right),
\end{aligned}$$

where $\mathrm{diag}(v_1,\ldots,v_l)$ denotes the diagonal matrix with entries $v_1,\ldots,v_l \in \mathbb{R}$ on its diagonal and δ is the Kronecker delta. In the second equation we made use of $E\left(N_{s,i}N_{u,a}N_{t,j}N_{v,b}\right) = \delta_{i,a}\delta_{j,b}\delta_{s,u}\delta_{t,v} + \delta_{i,b}\delta_{j,a}\delta_{s,v}\delta_{t,u} + \delta_{i,j}\delta_{a,b}\delta_{t,s}\delta_{u,v} + \kappa\,\delta_{t,s}\delta_{s,u}\delta_{u,v}\delta_{i,j}\delta_{j,a}\delta_{a,b}$. Hence,

$$\begin{aligned}
\frac{1}{\|V\|^4}E\left(q^{\mathrm{T}}A_0P^{\mathrm{T}}PA_0q\right) &= \frac{1}{\|V\|^4}\mathrm{tr}\left\{P^{\mathrm{T}}P\mathrm{var}\left(A_0q\right)\right\} - \frac{1}{\|V\|^4}E\left(q^{\mathrm{T}}A_0\right)P^{\mathrm{T}}PE\left(A_0q\right) \\
&= \frac{\|V^2\|^2}{\|V\|^4}\left(q^{\mathrm{T}}P^{\mathrm{T}}Pq + \|P\|^2\|q\|^2\right) + q^{\mathrm{T}}P^{\mathrm{T}}Pq + \kappa\sum_{t=1}^{n}\frac{\|V_t\|^4}{\|V\|^4}\sum_{i=1}^{l}\|P_i\|^2q_i^2.
\end{aligned}$$

The remaining term in S_1 follows trivially, proving the result. $E\|\Sigma^2 - A\|^2$ is obtained using similar calculations. □

S1·3. *Proof of Theorem 2*

LEMMA S1. *Assume that for $\nu \in (0,1]$ and some constants $\delta, \epsilon > 0$ it holds that* $\mathrm{pr}\left(\|A - \Sigma^2\|_{\mathcal{L}} \le \delta\right) \ge 1 - \nu/2$ *and* $\mathrm{pr}\left(\|b - Pq\| \le \epsilon\right) \ge 1 - \nu/2$. *Then each of the inequalities*

$$\begin{aligned}
\|A^{1/2} - \Sigma\| &\le 2^{-1}\delta\|\Sigma^{-1}\|\{1 + o(1)\}, \\
\|A^{-1/2}b - \Sigma^{-1}Pq\| &\le \epsilon\|\Sigma^{-1}\|_{\mathcal{L}} + 2^{-1}\delta(\|Pq\| + \epsilon)\|\Sigma^{-2}\|\|\Sigma^{-1}\|\{1 + o(1)\}
\end{aligned}$$

hold with probability at least $1 - \nu/2$.

Proof: We show the result by using the Fréchet-derivative for functions $F : \mathbb{R}^{k\times k} \to \mathbb{R}^{k\times k}$. Due to the fact that $\eta_1 > 0$ it holds that Σ^2 is positive definite and thus invertible.

It holds due to Higham (2008), Problem 7.4, that $F'(\Sigma^2)B$ for an arbitrary $B \in \mathbb{R}^{k\times k}$ is given as the solution in $X \in \mathbb{R}^{k\times k}$ of $B = \Sigma X + X\Sigma$, i.e. due to the symmetry and positive definitiness of Σ we have $F'(\Sigma^2)B = 2^{-1}\Sigma^{-1}B$. We take the orthonormal base $\{E_{i,j}, i,j = 1,\ldots,k\}$

for the space $(\mathbb{R}^{k\times k}, \|\cdot\|)$ with $E_{i,j}$ corresponding to the matrix that has zeros everywhere except at the position (i,j), where it is one. The Hilbert-Schmidt norm $\|F'(\Sigma^2)\|_{HS}$ is

$$\|F'(\Sigma^2)\|_{HS}^2 = 4^{-1}\sum_{i,j=1}^{k}\|\Sigma^{-1}E_{i,j}\|^2 = 4^{-1}\sum_{i,j=1}^{k}[\Sigma^{-1}]_{i,j}^2 = 4^{-1}\|\Sigma^{-1}\|^2.$$

This yields with the Taylor expansion for Fréchet-differentiable maps

$$\|A^{1/2}-\Sigma\|_{\mathcal{L}} \le \|F'(\Sigma)(A-\Sigma^2)\| + o(\|A-\Sigma^2\|) \le 2^{-1}\|\Sigma^{-1}\|\delta\{1+o(1)\}.$$

For the second inequality we see first that

$$\|A^{-1/2}b - \Sigma^{-1}Pq\| \le \epsilon\|\Sigma^{-1}\|_{\mathcal{L}} + \left\|(A^{-1/2}-\Sigma^{-1})b\right\|. \tag{S1}$$

The Fréchet-derivative of the map $F:\mathbb{R}^{k\times k}\to\mathbb{R}^{k\times k}, A\mapsto A^{-1/2}$ is $F'(\Sigma^2)B = -2^{-1}\Sigma^{-2}B\Sigma^{-1}$ and

$$\|F'(\Sigma^2)\|_{HS}^2 = 4^{-1}\sum_{i,j=1}^{k}\|\Sigma^{-2}E_{i,j}\Sigma^{-1}\|^2 \le 4^{-1}\|\Sigma^{-2}\|^2\|\Sigma^{-1}\|^2.$$

Here we used the submultiplicativity of the Frobenius norm with the Hadamard product of matrices. Thus we get via Taylor's theorem

$$\|A^{-1/2}-\Sigma^{-1}\| \le 2^{-1}\|\Sigma^{-2}\|\|\Sigma^{-1}\|\|A-\Sigma^2\| + o(\delta).$$

Plugging this into (S1) yields

$$\|A^{-1/2}b - \Sigma^{-1}Pq\| \le \epsilon\|\Sigma^{-1}\|_{\mathcal{L}} + 2^{-1}\delta(\|Pq\|+\epsilon)\|\Sigma^{-2}\|\|\Sigma^{-1}\|\left\{1+o(1)\right\},$$

where we used that $\|b\| \le \|Pq\| + \epsilon$. □

Equivalence of conjugate gradient and partial least squares: We denote $\tilde{A} = A^{1/2}$ and $\tilde{b} = A^{-1/2}b$. The partial least squares optimization problem is

$$\min_{v\in\mathcal{K}_i(A,b)}\|y-Xv\|^2,$$

whereas the conjugate gradient problem studied in Nemirovskii (1986) is

$$\min_{v\in\mathcal{K}_i(\tilde{A}^2,\tilde{A}\tilde{b})}\|\tilde{b}-\tilde{A}v\|^2. \tag{S2}$$

It is easy to see that the Krylov space $\mathcal{K}_i(\tilde{A}^2,\tilde{A}\tilde{b}) = \mathcal{K}_i(A,b)$ $(i=1,\dots,k)$. We have

$$\arg\min_{v\in\mathcal{K}_i(\tilde{A}^2,\tilde{A}\tilde{b})}\|\tilde{b}-\tilde{A}v\|^2 = \arg\min_{\mathcal{K}_i(A,b)}\|y-Xv\|^2, i=1,\dots,k.$$

Thus it holds

$$\widehat{\beta}_i = \arg\min_{v\in\mathcal{K}_i(\tilde{A}^2,\tilde{A}\tilde{b})}\|\tilde{b}-\tilde{A}v\|^2,$$

Furthermore we have $\Sigma\beta(\eta_1) = \Sigma^{-1}Pq$, i.e. the correct problem in the population is solved by $\beta(\eta_1)$ as well. Now we will restate the main result in Nemirovskii (1986) in our context:

THEOREM S1. *Nemirovskii*
Assume that there are $\tilde{\delta} = \tilde{\delta}(\nu,n) > 0$, $\tilde{\epsilon} = \tilde{\epsilon}(\nu,n) > 0$ *such that for* $\nu \in (0,1]$ *it holds that*

$\mathrm{pr}\left(\|\Sigma - A^{1/2}\|_{\mathcal{L}} \leq \tilde{\delta}\right) \geq 1-\nu/2$, $\mathrm{pr}\left(\|\Sigma^{-1}Pq - A^{-1/2}b\| \leq \tilde{\epsilon}\right) \geq 1-\nu/2$ *and the conditions*

1. *there is an* $L = L(\nu, n)$ *such that with probability at least* $1-\nu/2$ *it holds that* $\max\left\{\|A^{1/2}\|_{\mathcal{L}}, \|\Sigma\|_{\mathcal{L}}\right\} \leq L$*,*
2. *there is a vector* $u \in \mathbb{R}^k$ *and constants* $R, \mu > 0$ *such that* $\beta(\eta_1) = \Sigma^{\mu} u$*,* $\|u\| \leq R$

are satisfied. If we stop according to the stopping rule a^* *as defined in (4) with* $\tau \geq 1$ *and* $\zeta < \tau^{-1}$ *then we have for any* $\theta \in [0,1]$ *with probability at least* $1-\nu$

$$\left\|\Sigma^{\theta}\{\widehat{\beta}_{a^*} - \beta(\eta_1)\}\right\|^2 \leq C^2(\mu,\tau,\zeta) R^{2(1-\theta)/(1+\mu)} \left(\tilde{\epsilon} + \tilde{\delta} R L^{\mu}\right)^{2(\theta+\mu)/(1+\mu)}.$$

Proof: Note first that on the set where $\|\Sigma - A^{1/2}\|_{\mathcal{L}} \leq \tilde{\delta}$ holds with probability at least $1-\nu/2$ condition 1 also holds with $L = \|\Sigma\|_{\mathcal{L}} + \tilde{\delta}$. Constrained on the set where all the conditions of the theorem hold with probability at least $1-\nu$ we consider Nemirovskii's $(\Sigma, A^{1/2}, \Sigma^{-1}Pq, A^{-1/2}b)$ problem with errors $\tilde{\delta}$ and $\tilde{\epsilon}$. Furthermore by assumption Nemirovskii's $(2\theta, R, L, 1)$ conditions hold and thus the theorem follows by a simple application of the main theorem in Nemirovskii (1986). □

We will now apply Theorem S1 to our problem. Due to the fact that $\eta_1 > 0$ it holds that Σ^2 is positive definite and thus invertible. We note that the spectral norm is dominated by the Frobenius norm. From Markov's inequality we get

$$\mathrm{pr}\left(\|A - \Sigma^2\| \geq \delta\right) \leq \delta^{-2} E\left(\left\|A - \Sigma^2\right\|^2\right).$$

Using Theorem 1, $\sum_{t=1}^{n} \|V_i\|^4 \leq \|V^2\|^2$ and setting the right hand side to $\nu/2$ for $\nu \in (0,1]$ gives $\delta = \nu^{-1/2}\|V\|^{-2}\|V^2\|C_{\delta}$. In the same way $\epsilon = \nu^{-1/2}\|V\|^{-2}\|V^2\|C_{\epsilon}$. Lemma S1 gives with probability at least $1-\nu/2$ the concentration results required by Theorem S1 with

$$\begin{aligned}
\tilde{\delta} &= \nu^{-1/2}\frac{\|V^2\|}{\|V\|^2} C_{\delta}\{1+o(1)\} \\
\tilde{\epsilon} &= \left(\nu^{-1/2}\frac{\|V^2\|}{\|V\|^2} C_{\epsilon} + \nu^{-1}\frac{\|V^2\|^2}{\|V\|^4} C_{\epsilon} C_{\delta}\right)\{1+o(1)\}
\end{aligned}$$

Conditions 1 and 2 of Theorem S1 hold with a probability of at least $1-\nu/2$ by choosing $L = \tilde{\delta} + \|\Sigma\|_{\mathcal{L}}$, $\mu = 1$ and $R = \|\Sigma^{-3}Pq\|$. Here we used that $\beta(\eta_1) = \Sigma^{-2}Pq$. Thus the theorem yields for $\theta = 1$

$$\left\|\Sigma\{\beta(\eta_1) - \widehat{\beta}_{a^*}\}\right\| \leq C(1,\tau,\zeta)\left(\tilde{\epsilon} + \tilde{\delta} R L\right).$$

Denote $c(\tau,\zeta) = C(1,\tau,\zeta)\{1+o(1)\}$. Finally we have $\|\Sigma^{-1}\|_{\mathcal{L}}^{-1}\|v\| \leq \|\Sigma v\|$ for any $v \in \mathbb{R}^k$ and thus the theorem is proven with

$$\begin{aligned}
c_1(\nu) &= \nu^{-1/2} c(\tau,\zeta)\|\Sigma^{-1}\|_{\mathcal{L}}\left(C_{\epsilon} + \|\Sigma\|_{\mathcal{L}}\|\Sigma^{-3}Pq\|C_{\delta}\right) \\
c_2(\nu) &= \nu^{-1} c(\tau,\zeta)\|\Sigma^{-1}\|_{\mathcal{L}}\left(C_{\epsilon}C_{\delta} + \|\Sigma^{-3}Pq\|C_{\delta}^2\right).
\end{aligned}$$

□

S1·4. *Proof of Theorem 3*

The theorem is proved by contradiction. Assume that $\widehat{\beta}_1 \longrightarrow \beta_1$ in probability. Choosing $v \in \mathbb{R}^k$, $v \neq 0$, orthogonal to β_1 implies that $v^{\mathrm{T}}\widehat{\beta}_1$ converges in probability to zero. Next we show that the second moment vanishes as well.

Let $M_d(z) = \max_{i\in\{1,\dots,n\}^d} E(\prod_{\nu=1}^d z_{i_\nu}^2)$ for a random vector $z = (z_1, \dots, z_n)^{\mathrm{T}}$ with existing mixed $(2d)$th moments. Using $(a+b)^2 \leq 2(a^2+b^2)$ for $a, b \in \mathbb{R}$ we obtain

$$\begin{aligned} E\left(v^{\mathrm{T}}b\right)^4 \leq & \frac{8^2\|v\|^4}{\|V\|^8} E\left(\left\|PN^{\mathrm{T}}V^2Nq\right\|^4 + \eta_1^4\left\|F^{\mathrm{T}}V^2Nq\right\|^4 + \eta_2^4\left\|PN^{\mathrm{T}}V^2f\right\|^4 + \eta_1^4\eta_2^4\left\|F^{\mathrm{T}}V^2f\right\|^4\right) \\ \leq & 8^2\|v\|^4\left\{M_4(N_1)\|q\|^4l^4\|P\|^4 + M_2(N_1)M_2(F_1)\eta_1^4\|q\|^4l^2k^2\right. \\ & \left. + M_2(N_1)M_2(f_1)\eta_2^4l^2\|P\|^4 + M_2(F_1)M_2(f_1)\eta_1^4\eta_2^4k^2\right\} < \infty, \;\; n \in \mathbb{N}. \end{aligned}$$

Thus, $(v^{\mathrm{T}}b)^2$ is uniformally integrable by the theorem of de la Vallée-Poussin and it follows that the directional variance $\mathrm{var}(v^{\mathrm{T}}b)$ has to vanish in the limit as well. Now, calculations similar to Theorem 1 yield

$$\begin{aligned} \mathrm{var}(v^{\mathrm{T}}b) = & \frac{\|V^2\|^2}{\|V\|^4}\left\{\eta_1^2\|v\|^2\left(\|q\|^2+\eta_2^2\right) + \|P^{\mathrm{T}}v\|^2\left(\|q\|^2+\eta_2^2\right) + (v^{\mathrm{T}}Pq)^2\right\} \\ & + \sum_{t=1}^n \frac{\|V_t\|^4}{\|V\|^4}\sum_{i=1}^l q_i^2\left(v^{\mathrm{T}}P_i\right)^2\left\{E(N_{1,1}^4) - 3\right\}, \;\; v \in \mathbb{R}^k. \end{aligned}$$

We assumed that $\|V\|^{-2}\|V^2\|$ does not converge to zero. It remains to check under which conditions $\mathrm{var}(v^{\mathrm{T}}b)$ is larger than zero. This will always be the case if $v \neq 0$ and $\eta_1 > 0$, $l = 1$. For $\eta_1 = 0$ and $l > 1$ a vector v that lies in the range of P and is orthogonal to $\beta_1 \propto Pq$ exists, thus contradicting $\widehat{\beta}_1 \longrightarrow \beta_1$ in probability. □

S1·5. *Proof of Theorem 4*

It is easy to verify that $\|V\|^2 = \mathrm{tr}(T^2) = n\gamma(0)$ and $\left\|V^2\right\|^2 = n\gamma^2(0) + 2\sum_{t=1}^{n-1}\gamma^2(t)(n-t)$. If **(??)** is fulfilled, then

$$n\gamma(0) \leq \left\|V^2\right\|^2 \leq n\gamma^2(0)\left\{1 + 2c^2\frac{1-\exp(-2\rho(n-1))}{\exp(2\rho)-1}\right\} \leq n\gamma^2(0)\left\{1 + \frac{2c}{\exp(2\rho)-1}\right\}.$$

It follows that $\|V^2\| \sim n^{1/2}$. □

S1·6. *Proof of Theorem 5*

Let $\gamma : \mathbb{N} \to \mathbb{R}$ be the autocovariance function of a stationary time series that has zero mean. For the autocovariance matrix V^2 of the corresponding integrated process of order one we get $\left[V^2\right]_{t,s} = \sum_{i,j=1}^{t,s}\gamma(|i-j|)$, $(t, s = 1, \dots, n)$. Let $t \geq s$. By splitting the sum into parts with $i < j$ and $i > j$ we get $\left[V^2\right]_{t,s} = s\gamma(0) + \sum_{j=1}^s\sum_{i=1}^{t-j}\gamma(i) + \sum_{j=2}^s\sum_{i=1}^{j-1}\gamma(i)$. Due to symmetry, $\left[V^2\right]_{t,s} = \left[V^2\right]_{s,t}$ for $s > t$.

First, consider the case that all $\gamma(j)$, $j > 0$ are negative. Using (6) we obtain

$$\gamma(0)s \geq \left[V^2\right]_{t,s} \geq \gamma(0)\left\{s - c\sum_{j=1}^s\sum_{i=1}^{t-j}\exp(-\rho j) - c\sum_{j=2}^s\sum_{i=1}^{j-1}\exp(-\rho j)\right\}, \; t \geq s.$$

Evaluation of the geometric sums gives

$$\left[V^2\right]_{t,s} \geq \gamma(0)\left(s\left\{1-\frac{2c}{\exp(\rho)-1}\right\}+c\frac{\exp(\rho)}{\{\exp(\rho)-1\}^2}\left\{1-\exp(-\rho s)\right\}\left[1+\exp\{\rho(s-t)\}\right]\right).$$

The second term on the right is always positive and the positivity of the first term is ensured by the condition $\rho > \log(2c+1)$. Hence, $\gamma(0)\left[1-2c\left\{\exp(\rho)-1\right\}^{-1}\right]s \leq \left[V^2\right]_{t,s} \leq \gamma(0)s,\ s \geq 1$. If $\gamma(t)$, $t \geq 1$ is not purely negative, it can be bound by

$$\gamma(0)\left[1-2c\left\{\exp(\rho)-1\right\}^{-1}\right]s \leq \left[V^2\right]_{t,s} \leq \gamma(0)\left[1+2c\left\{\exp(\rho)-1\right\}^{-1}\right]s.$$

We write δ_1 and δ_2 for the constants in the lower and upper bound, respectively, so that $\delta_1 \min\{s,t\} \leq \left[V^2\right]_{t,s} \leq \delta_2 \min\{s,t\}$ $(t,s=1,\ldots,n)$. This yields upper and lower bounds on the trace of V^2 and shows that $\|V\|^2 \sim n^2$. Additionally,

$$\left[V^4\right]_{t,t} = \sum_{l=1}^{n}\left[V^2\right]_{t,l}\left[V^2\right]_{l,t} = \sum_{l=1}^{t}\left[V^2\right]_{t,l}^2 + \sum_{l=t+1}^{n}\left[V^2\right]_{l,t}^2 \leq \frac{\delta_2^2}{6}t\left(6nt-4t^2+3t+1\right)$$

$$\left[V^4\right]_{t,t} \geq \frac{\delta_1^2}{6}t\left(6nt-4t^2+3t+1\right).$$

This implies upper and lower bounds on the trace of V^4 in the form $c\,n(n+1)(n^2+n+1)$ for $c \in \{\delta_1^2/6, \delta_2^2/6\}$ and thus $\|V^2\| \sim n^2$. □

S1·7. *Proof of Theorem 6*

First consider $n^{-1}X^{\mathrm{T}}\widehat{V}^{-2}y$. Define $X_u = (X_{u,1},\ldots,X_{u,n})^{\mathrm{T}} = NP^{\mathrm{T}} + \eta_1 F$ and $y_u = (y_{u,1},\ldots,y_{u,n})^{\mathrm{T}} = Nq + \eta_2 f$ such that $X = VX_u$ and $y = Vy_u$. By the triangle inequality

$$\left\|n^{-1}X^{\mathrm{T}}\widehat{V}^{-2}y - Pq\right\| \leq \left\|n^{-1}X^{\mathrm{T}}V^{-2}y - Pq\right\| + \left\|n^{-1}X^{\mathrm{T}}\left(\widehat{V}^{-2}-V^{-2}\right)y\right\|.$$

The first term on the right hand side is convergent to zero in probability due to Theorem 1. The second term can be bounded as

$$n^{-2}\left\|X^{\mathrm{T}}\left(\widehat{V}^{-2}-V^{-2}\right)y\right\|^2 \leq \|V\widehat{V}^{-2}V - I_n\|_{\mathcal{L}}^2\, n^{-1}\|X_u^{\mathrm{T}}\|_{\mathcal{L}}^2\, n^{-1}\|y_u\|^2.$$

Since both $X_{u,1},\ldots,X_{u,n}$ and $y_{u,1},\ldots,y_{u,n}$ are independent and identically distributed, it follows that $n^{-1}\|y_u\|^2$ is a strongly consistent estimator for $E(y_{u,1}^2)$, as well as that $n^{-1}\|X_u^{\mathrm{T}}\|_{\mathcal{L}}^2$ is bounded from above by $n^{-1}\|X_u^{\mathrm{T}}\|^2$, which is a strongly consistent estimator of $E(\|X_{u,1}\|^2)$. Convergence in probability of $\left\|V\widehat{V}^{-2}V - I_n\right\|_{\mathcal{L}}^2$ to zero implies the convergence of $b(\widehat{V})$ to Pq in probability. To obtain the convergence rate $\|n^{-1}X^{\mathrm{T}}V^{-2}y - Pq\| = O_p(r_n)$, use Theorem 1 and $\|V\widehat{V}^{-2}V - I_n\|_{\mathcal{L}} = O_p(r_n)$. The convergence of $\|n^{-1}X^{\mathrm{T}}\widehat{V}^{-2}X - \Sigma^2\|$ is proven in a similar way.

To show the consistency and the rate of the corrected partial least squares estimator, we follow the same lines as in the proof of Theorem 2. First, $\delta = r_n c_A(\nu)$ and $\epsilon = r_n c_b(\nu)$ for $\nu \in (0,1]$ with constants $c_A(\nu)$, $c_b(\nu)$ are taken, such that

$$\mathrm{pr}\{\|A(\widehat{V})^{1/2} - \Sigma\|_{\mathcal{L}} \leq r_n c_A(\nu)\} \geq 1-\nu/2,$$
$$\mathrm{pr}\{\|A(\widehat{V})^{-1/2}b(\widehat{V}) - \Sigma^{-1}Pq\| \leq r_n c_b(\nu)\} \geq 1-\nu/2.$$

Moreover, $L = \|\Sigma\|_{\mathcal{L}} + \delta$ and $R = \|\Sigma^{-3}Pq\|$, $\mu = 1$, satisfies conditions 1 and 2 in Theorem S1 with probability at least $1 - \nu/2$. Thus, with probability at least $1 - \nu$ we get by setting $\theta = 1$

$$\left\|\widehat{\beta}_{a^*}(\widehat{V}) - \beta(\eta_1)\right\| \leq r_n\, C(1,\tau,\zeta)\{1 + o(1)\}\|\Sigma^{-1}\|_{\mathcal{L}}\left[c_b(\nu) + c_A(\nu)\|\Sigma^{-3}Pq\|\left\{\|\Sigma\|_{\mathcal{L}} + r_n c_A(\nu)\right\}\right],$$

where the constants ζ, τ are taken from the definition of a^*. □

S2. Addendum to section 5, Simulations

Figure S1 shows the differences in empirical mean squared error of $\hat{\beta}_1$ for various dependence structures considered in Section 5 in the setting with $l = i = 1$. We calculated

$$n\text{MSE}(\hat{\beta}_1) = n\, 500^{-1} \sum_{i=1}^{500} (\hat{\beta}_{1,i} - \beta_1)^2,$$

where $\hat{\beta}_{1,i}$ denotes a partial least squares estimator in the ith Monte Carlo simulation based on n observations. If an autoregressive dependence is present in the data and is ignored in the partial least squares algorithm, $n\text{MSE}(\hat{\beta}_1)$ is proportional to a constant, which is larger than in the corrected partial least squares case. Ignoring the integrated dependence in the data leads to $n\text{MSE}(\hat{\beta}_1)$ growing linearly with n, which confirms our theoretical findings in Section 3.

References

HIGHAM, N. (2008). *Functions of Matrices: Theory and Computation*. Phildadelphia: SIAM, 1st ed.

NEMIROVSKII, A. (1986). The regularizing properties of the adjoint gradient method in ill-posed problems. *Comput. Math. Math. Phys.* **26**, 7–16.

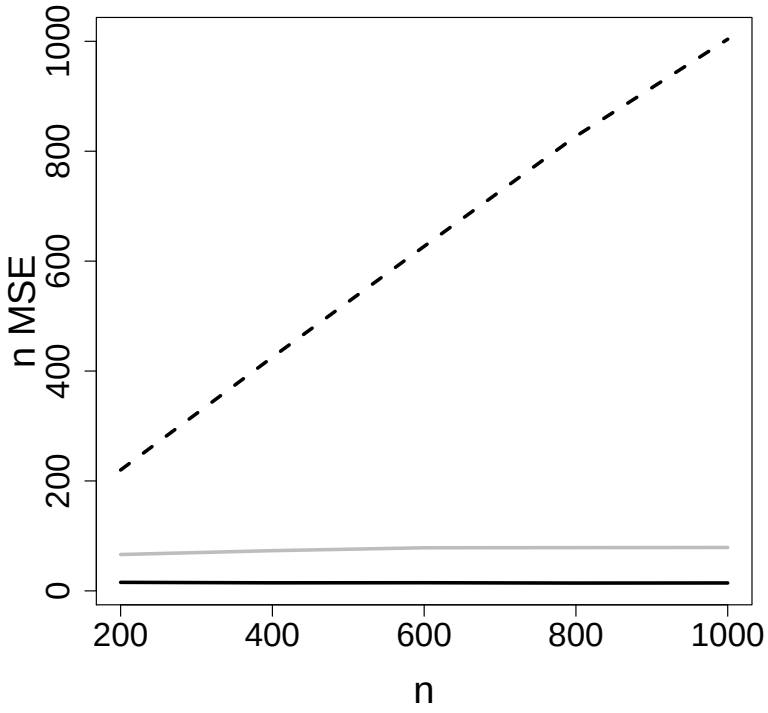

Fig. S1: Empirical mean squared eror of $\widehat{\beta}_1$ multiplied by n. The dependence structures are: autoregressive (grey), autoregressive integrated moving average (black, dashed) and corrected partial least squares on integrated data (black, solid).